The Axiomatic Method in Social Choice Theory:

Preference Aggregation, Judgment Aggregation, Graph Aggregation

Ulle Endriss
Institute for Logic, Language and Computation
University of Amsterdam

Social Choice Theory

SCT studies collective decision making: how should we aggregate the preferences of the members of a group to obtain a "social preference"?

Expert 1: $\triangle \succ \bigcirc \succ \square$

Expert 2: $\bigcirc \succ \square \succ \triangle$

Expert 3: $\square \succ \triangle \succ \bigcirc$

Expert 4: $\square \succ \triangle \succ \bigcirc$

Expert 5: $\bigcirc \succ \square \succ \triangle$

?

Outline

This will be an introduction to the *axiomatic method* in SCT:

- preference aggregation
- judgment aggregation
- graph aggregation

Background reading on PA and JA: see expository papers cited below.
The material on GA is based on original work with Umberto Grandi.

U. Endriss. Logic and Social Choice Theory. In A. Gupta and J. van Benthem (eds.), *Logic and Philosophy Today*, College Publications, 2011.

U. Endriss. Judgment Aggregation. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A.D. Procaccia (eds.), *Handbook of Computational Social Choice*. CUP, 2015.

U. Endriss and U. Grandi. Collective Rationality in Graph Aggregation. Proc. 21st European Conference on Artificial Intelligence (ECAI-2014).

Framework 1: Preference Aggregation

Basic terminology and notation:

- finite set of *individuals* $\mathcal{N} = \{1, \ldots, n\}$, with $n \geqslant 2$ *odd*
- (usually finite) set of *alternatives* $\mathcal{X} = \{x_1, x_2, x_3, \ldots\}$
- Denote the set of *linear orders* on $\mathcal{X}$ by $\mathcal{L}(\mathcal{X})$.
  Preferences (or *ballots*) are taken to be elements of $\mathcal{L}(\mathcal{X})$.
- A *profile* $\boldsymbol{R} = (R_1, \ldots, R_n) \in \mathcal{L}(\mathcal{X})^n$ is a vector of preferences.
- We shall write $N^{\boldsymbol{R}}_{x \succ y}$ for the set of individuals that rank alternative x above alternative y under profile $\boldsymbol{R}$.

We are interested in preference aggregation methods that map any profile of preferences to a single collective preference.

The proper technical term is *social welfare function* (SWF):

$$F : \mathcal{L}(\mathcal{X})^n \to \mathcal{L}(\mathcal{X})$$

Three Axioms

Axioms in SCT are mathematically rigorous encodings of normative requirements on aggregation methods. Three examples:

- F is *anonymous* if $F(R_1, \ldots, R_n) = F(R_{\pi(1)}, \ldots, R_{\pi(n)})$ for any profile $(R_1, \ldots, R_n)$ and any permutation $\pi : \mathcal{N} \to \mathcal{N}$.
- F is *neutral* if $F(\pi(\boldsymbol{R})) = \pi(F(\boldsymbol{R}))$ for any profile $\boldsymbol{R}$ and any permutation $\pi : \mathcal{X} \to \mathcal{X}$ (extended to preferences and profiles).
- F is (weakly) *monotonic* if, whenever $x \succ y$ in the outcome, then one additional agent adopting $x \succ y$ does not change this.

May's Theorem

Example for a characterisation result (useful to justify a rule):

Theorem 1 (May, 1952) *In case of two alternatives, a rule is anonymous, neutral, and monotonic iff it is the simple majority rule.*

Proof: ($\Leftarrow$) Obvious. ✓ ($\Rightarrow$) Everyone votes either $x \succ y$ or $y \succ x$. *ANON* $\rightsquigarrow$ only number of ballots of each type matters. Two cases:

- Suppose $|N^{\boldsymbol{R}}_{x \succ y}| = |N^{\boldsymbol{R}}_{y \succ x}| + 1$ implies $(x \succ y) = F(\boldsymbol{R})$. Then, by *MONO*, F must be the simple majority rule. ✓
- Suppose $\exists \boldsymbol{R}$ s.t. $|N^{\boldsymbol{R}}_{x \succ y}| = |N^{\boldsymbol{R}}_{y \succ x}| + 1$ but $(y \succ x) = F(\boldsymbol{R})$. Let one voter switch from $x \succ y$ to $y \succ x$ to yield $\boldsymbol{R'}$. Then by *NEUT* $(x \succ y) = F(\boldsymbol{R'})$, but by *MONO* $(y \succ x) = F(\boldsymbol{R'})$. ↯

Note: This result is usually presented in a slightly different framework.

K.O. May. A Set of Independent Necessary and Sufficient Conditions for Simple Majority Decisions. *Econometrica*, 20(4):680–684, 1952.

Two More Axioms

Back to the case of arbitrary numbers of alternatives . . .

- F satisfies the (weak) *Pareto condition* if, whenever all individuals rank x above y, then so does society:

$$N^{\boldsymbol{R}}_{x \succ y} = \mathcal{N} \text{ implies } (x \succ y) \in F(\boldsymbol{R})$$

- F satisfies *independence of irrelevant alternatives* (*IIA*) if the relative social ranking of two alternatives only depends on their relative individual rankings:

$$N^{\boldsymbol{R}}_{x \succ y} = N^{\boldsymbol{R'}}_{x \succ y} \text{ implies } (x \succ y) \in F(\boldsymbol{R}) \Leftrightarrow (x \succ y) \in F(\boldsymbol{R'})$$

  In other words: if x is socially preferred to y, then this should not change when an individual changes her ranking of z.

Arrow's Theorem

A SWF F is a *dictatorship* if there exists a "dictator" $i \in \mathcal{N}$ such that $F(\boldsymbol{R}) = R_i$ for any profile $\boldsymbol{R}$, i.e., if the outcome is always identical to the preference supplied by the dictator.

Theorem 2 (Arrow, 1951) *Any SWF for $\geqslant 3$ alternatives that satisfies the Pareto condition and IIA must be a dictatorship.*

Proof: Omitted (more difficult than for May's Theorem).

Remarks:

- surprising / not true for 2 alternatives / opposite direction clear
- dictatorship does not just mean "someone agrees with outcome"
- impossibility result = characterisation of bad SWF (dictatorship)
- historical significance: message / generality / methodology

K.J. Arrow. *Social Choice and Individual Values*. John Wiley and Sons, 2nd edition, 1963. First edition published in 1951.

Example: Judgment Aggregation

	p	$p \to q$	q
Judge 1:	True	True	True
Judge 2:	True	False	False
Judge 3:	False	True	False

?

Framework 2: Judgment Aggregation

Notation: Let $\sim\varphi := \varphi'$ if $\varphi = \neg\varphi'$ and let $\sim\varphi := \neg\varphi$ otherwise.

An *agenda* Φ is a finite nonempty set of propositional formulas (w/o double negation) closed under complementation: $\varphi \in \Phi \;\Rightarrow\; \sim\varphi \in \Phi$.

A *judgment set* J on an agenda Φ is a subset of Φ. We call J:

- *complete* if $\varphi \in J$ or $\sim\varphi \in J$ for all $\varphi \in \Phi$
- *complement-free* if $\varphi \notin J$ or $\sim\varphi \notin J$ for all $\varphi \in \Phi$
- *consistent* if there exists an assignment satisfying all $\varphi \in J$

Let $\mathcal{J}(\Phi)$ be the set of all complete and consistent subsets of Φ.

A finite set of *individuals* $\mathcal{N} = \{1, \ldots, n\}$, with $n \geqslant 2$ *odd*, express judgments on the formulas in Φ, producing a *profile* $\boldsymbol{J} = (J_1, \ldots, J_n)$.

An *aggregation rule* for an agenda Φ and a set of n individuals is a function mapping a profile of complete and consistent individual judgment sets to a single collective judgment set: $F : \mathcal{J}(\Phi)^n \to 2^\Phi$.

Example: Majority Rule

Suppose three agents ($\mathcal{N} = \{1, 2, 3\}$) express judgments on the propositions in the agenda $\Phi = \{p, \neg p, q, \neg q, p \vee q, \neg(p \vee q)\}$.

For simplicity, we only show the positive formulas in our tables:

	p	q	$p \vee q$	
Agent 1:	True	False	True	$J_1 = \{p, \neg q, p \vee q\}$
Agent 2:	True	True	True	$J_2 = \{p, q, p \vee q\}$
Agent 3:	False	False	False	$J_3 = \{\neg p, \neg q, \neg(p \vee q)\}$

The (strict) *majority rule* $F_{\textsf{maj}}$ takes a (complete and consistent) profile and returns the set of propositions accepted by $> \frac{n}{2}$ agents.

In our example: $F_{\textsf{maj}}(\boldsymbol{J}) = \{p, \neg q, p \vee q\}$ [complete and consistent!]

In general, $F_{\textsf{maj}}$ only ensures completeness and complement-freeness [and completeness only in case n is odd].

Some Axioms

What makes for a "good" aggregation rule F? The following *axioms* all express intuitively appealing (yet, debatable) properties:

- *Anonymity:* Treat all individuals symmetrically!

  Formally: for any profile $\boldsymbol{J}$ and any permutation $\pi : \mathcal{N} \to \mathcal{N}$ we have $F(J_1, \ldots, J_n) = F(J_{\pi(1)}, \ldots, J_{\pi(n)})$.

- *Neutrality:* Treat all propositions symmetrically!

  Formally: for any φ, ψ in the agenda Φ and any profile $\boldsymbol{J}$, if for all $i \in \mathcal{N}$ we have $\varphi \in J_i \Leftrightarrow \psi \in J_i$, then $\varphi \in F(\boldsymbol{J}) \Leftrightarrow \psi \in F(\boldsymbol{J})$.

- *Independence:* Only the "pattern of acceptance" should matter!

  Formally: for any φ in the agenda Φ and any profiles $\boldsymbol{J}$ and $\boldsymbol{J'}$, if $\varphi \in J_i \Leftrightarrow \varphi \in J'_i$ for all $i \in \mathcal{N}$, then $\varphi \in F(\boldsymbol{J}) \Leftrightarrow \varphi \in F(\boldsymbol{J'})$.

Observe that the *majority rule* satisfies all of these axioms.

(But so do some other procedures! Can you think of some examples?)

Impossibility Theorem

We have seen that the majority rule is *not consistent*. Is there some other "reasonable" aggregation rule that does not have this problem? *Surprisingly, no!* (at least not for certain agendas)

Theorem 3 (List and Pettit, 2002) *No judgment aggregation rule for an agenda Φ with $\{p, q, p \wedge q\} \subseteq \Phi$ that satisfies the axioms of anonymity, neutrality, and independence will always return a collective judgment set that is complete and consistent.*

Remark 1: Note that the theorem requires $|\mathcal{N}| > 1$.

Remark 2: Similar impossibilities arise for other agendas with some minimal structural richness.

C. List and P. Pettit. Aggregating Sets of Judgments: An Impossibility Result. *Economics and Philosophy*, 18(1):89–110, 2002.

Proof

Let $N_\varphi^{\boldsymbol{J}}$ be the set of individuals who accept formula φ in profile $\boldsymbol{J}$.

Let F be any aggregator that is independent, anonymous, and neutral.

- Due to *IND*, whether $\varphi \in F(\boldsymbol{J})$ only depends on $N_\varphi^{\boldsymbol{J}}$.
- Then, by *ANON*, whether $\varphi \in F(\boldsymbol{J})$ only depends on $|N_\varphi^{\boldsymbol{J}}|$.
- But, by *NEUT*, how $\varphi \in F(\boldsymbol{J})$ depends on $|N_\varphi^{\boldsymbol{J}}|$ mustn't dep. on φ.

Thus: if φ and ψ are accepted by the same number of individuals, then we must either accept both of them or reject both of them.

Consider a profile $\boldsymbol{J}$ where $\frac{n-1}{2}$ individuals accept p and q; one accepts p but not q; one accepts q but not p; and $\frac{n-3}{2}$ accept neither p nor q. Thus: $|N_p^{\boldsymbol{J}}| = |N_q^{\boldsymbol{J}}| = |N_{\neg(p\wedge q)}^{\boldsymbol{J}}| = \frac{n+1}{2}$ (recall: n is odd). Then:

- Accepting all three formulas contradicts consistency. ✓
- But if we accept none, completeness forces us to accept their complements, which also contradicts consistency. ✓

Graph Aggregation

Judgment aggregation generalises preference aggregation: you can judge propositions such as "$x \succ y$". A middle way is graph aggregation.

Fix a finite set of *vertices* V. A (directed) *graph* $G = \langle V, E\rangle$ based on V is defined by a set of *edges* $E \subseteq V \times V$ (thus: graph = edge-set).

Everyone in a finite group of *agents* $\mathcal{N} = \{1, \ldots, n\}$ provides a graph, giving rise to a *profile* $\boldsymbol{E} = (E_1, \ldots, E_n)$.

An *aggregator* is a function mapping profiles to collective graphs:

$$F : (2^{V \times V})^n \to 2^{V \times V}$$

Example: *majority rule* (accept an edge *iff* $> \frac{n}{2}$ of the individuals do)

U. Endriss and U. Grandi. Collective Rationality in Graph Aggregation. Proc. 21st European Conference on Artificial Intelligence (ECAI-2014).

Axioms

We may want to impose certain *axioms* on $F : (2^{V \times V})^n \to 2^{V \times V}$, e.g.:

- *Anonymous:* $F(E_1, \ldots, E_n) = F(E_{\sigma(1)}, \ldots, E_{\sigma(n)})$
- *Nondictatorial:* for no $i^\star \in \mathcal{N}$ you always get $F(\boldsymbol{E}) = E_{i^\star}$
- *Unanimous:* $F(\boldsymbol{E}) \supseteq E_1 \cap \cdots \cap E_n$
- *Grounded:* $F(\boldsymbol{E}) \subseteq E_1 \cup \cdots \cup E_n$
- *Neutral:* $N_e^{\boldsymbol{E}} = N_{e'}^{\boldsymbol{E}}$ implies $e \in F(\boldsymbol{E}) \Leftrightarrow e' \in F(\boldsymbol{E})$
- *Independent:* $N_e^{\boldsymbol{E}} = N_e^{\boldsymbol{E'}}$ implies $e \in F(\boldsymbol{E}) \Leftrightarrow e \in F(\boldsymbol{E'})$

For technical reasons, we'll restrict some axioms to *nonreflexive edges* $(x, y) \in V \times V$ with $x \neq y$ (NR-neutral, NR-nondictatorial).

Notation: $N_e^{\boldsymbol{E}} = \{i \in \mathcal{N} \mid e \in E_i\}$ = *coalition* accepting edge e in $\boldsymbol{E}$

Collective Rationality

Aggregator F is *collectively rational* (CR) for graph property P if, whenever all individual graphs E_i satisfy P, so does the outcome $F(\boldsymbol{E})$.

Examples for *graph properties:* reflexivity, transitivity, seriality, ...

Example

Three agents each provide a graph on the same set of four vertices:

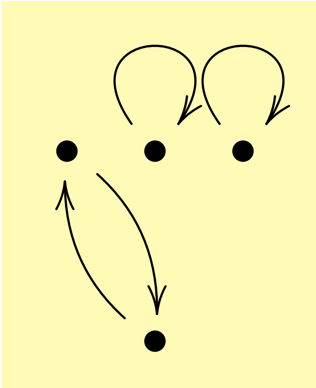

1

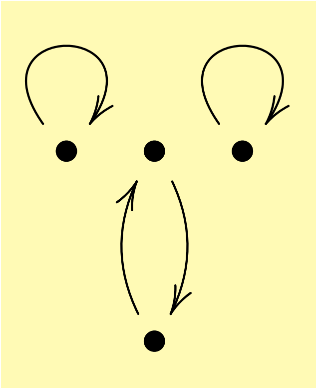

2

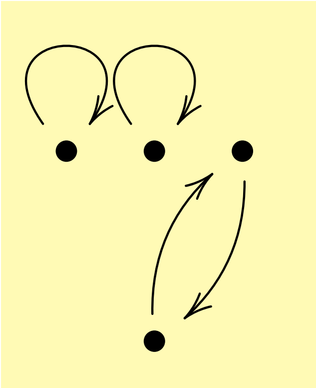

3

If we aggregate using the *majority rule*, we obtain this graph:

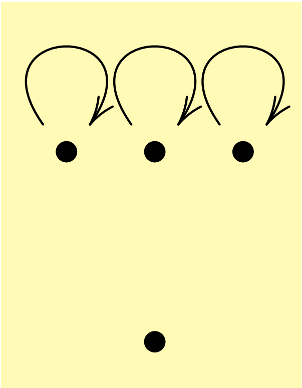

Observations:

- Majority rule not collectively rational for *seriality*.
- But *symmetry* is preserved.
- So is *reflexivity* (easy: individuals violate it).

Our General Impossibility Theorem

Our main result:

> *For $|V| \geqslant 3$, there exists no NR-nondictatorial, unanimous, grounded, and independent aggregator that is CR for any graph property that is contagious, implicative and disjunctive.*

where:

- Implicative $\approx [\bigwedge S^+ \wedge \neg \bigvee S^-] \to [e_1 \wedge e_2 \to e_3]$
- Disjunctive $\approx [\bigwedge S^+ \wedge \neg \bigvee S^-] \to [e_1 \vee e_2]$
- Contagious $\approx$ for every accepted edge, there are some conditions under which also one of its "neighbouring" edges is accepted

Examples:

- *Transitivity* is contagious and implicative
- *Completeness* is disjunctive

$\Big\}$ $\Rightarrow$ *Arrow's Theorem*

- *Connectedness* $[xEy \wedge xEz \to (yEz \vee zEy)]$ has all 3 properties

Last Slide

Social choice theory deals with the *aggregation of information* supplied by several individuals into a single such piece of information.

The traditional framework is that of *preference aggregation*, but other types of information (*judgments*, *graphs*, ...) are also of intereest.

The *axiomatic method* is maybe the most important classical method for studying aggregation—but there's much more to SCT/COMSOC.

We have seen:

- *axioms* such as anonymity, independence, monotonicity, ...
- *characterisation* (May) + *impossibility* (Arrow, List-Pettit) results
- a glimpse at *proof methods* for the simpler results
- a hint at the interplay of *axioms* with *collective rationality*